arXiv:1405.4628v1 [cs.IT] 19 May 2014

Distributed noise-shaping quantization: I. Beta duals of finite frames and near-optimal quantization of random measurements

Evan Chou* and C. Sinan Güntürk†
Courant Institute, NYU

May 20, 2014

Abstract

This paper introduces a new algorithm for the so-called "Analysis Problem" in quantization of finite frame representations which provides a near-optimal solution in the case of random measurements. The main contributions include the development of a general quantization framework called *distributed noise-shaping*, and in particular, *beta duals* of frames, as well as the performance analysis of beta duals in both deterministic and probabilistic settings. It is shown that for Gaussian random frames, using beta duals result in near-optimally accurate reconstructions with respect to both the frame redundancy and the number of levels that the frame coefficients are quantized at. More specifically, if L quantization levels per measurement are used to encode the unit ball in $\mathbb{R}^k$ via a Gaussian frame of m vectors, then with overwhelming probability the beta-dual reconstruction error is shown to be bounded by $\sqrt{k}L^{-(1-\eta)m/k}$ where η is arbitrarily small for sufficiently large problems. Additional features of the proposed algorithm include low computational cost and parallel implementability.

Keywords Finite frames, quantization, random matrices, noise shaping, beta encoding.

Mathematics Subject Classification 41A29, 94A29, 94A20, 42C15, 15B52

1 Introduction

Let $\mathcal{X}$ denote a bounded subset of $\mathbb{R}^k$, such as the unit ball B_1 with respect to the ℓ^2-norm, and $\mathcal{A}$ denote a quantization alphabet (a finite subset of $\mathbb{R}$), such as an arithmetic progression. Given a finite frame, i.e., a list of m vectors $F := \{f_1, \dots, f_m\}$ that span $\mathbb{R}^k$, we consider

$$\mathcal{D}_{\rm s}(F, \mathcal{A}, \mathcal{X}) := \sup_{x \in \mathcal{X}} \inf_{q \in \mathcal{A}^m} \left\| x - \sum_{i=1}^m q_i f_i \right\|_2 \tag{1}$$

which represents the smallest error of approximation that can be attained uniformly over $\mathcal{X}$ by considering linear combinations of elements of F with coefficients in $\mathcal{A}$. We will call this quantity the *synthesis distortion* associated with $(F, \mathcal{A}, \mathcal{X})$ (hence the subscript 's'). Understanding the

*Current affiliation: Google New York. email: chou@cims.nyu.edu
†email: gunturk@cims.nyu.edu

nature of the synthesis distortion and finding quantization algorithms (i.e., maps $x \in \mathcal{X} \mapsto q \in \mathcal{A}^m$) that achieve or even approximate this quantity is generally an open problem, usually referred to as the "Synthesis Problem" in the theory of quantization of frame representations [18]. The synthesis distortion can also be phrased as a generalized linear discrepancy [5, 14, 17].

Consider any other frame $E := \{e_1, \ldots, e_m\}$ in $\mathbb{R}^k$ that is dual to F, i.e.

$$\sum_{i=1}^{m} \langle x, e_i \rangle f_i = x \text{ for all } x \in \mathbb{R}^k. \tag{2}$$

In this paper we will adopt a matrix-based notation, and write F for the matrix whose columns are $(f_i)_1^m$ and E for the matrix whose *rows* are $(e_i^\top)_1^m$. (Depending on the role it assumes, E will be called the analysis frame or the sampling operator, and similarly F will be called the synthesis frame or the reconstruction operator.) It is then evident that (2) is equivalent to $FE = I$ and also that duality of frames is a symmetric relation. We define the *analysis distortion* associated with $(E, \mathcal{A}, \mathcal{X})$ by

$$\mathcal{D}_{\mathrm{a}}(E, \mathcal{A}, \mathcal{X}) := \inf\{\mathcal{D}_{\mathrm{s}}(F, \mathcal{A}, \mathcal{X}) : FE = I\}. \tag{3}$$

The analysis distortion has the following interpretation: Given an analysis frame E, the encoder chooses a quantization map $Q : \mathbb{R}^m \to \mathcal{A}^m$ and the decoder chooses a synthesis frame F that is dual to E. The encoder maps each $x \in \mathcal{X}$ (or equivalently, the measurement vector $y := Ex$) to its quantized version $q := Q(Ex)$), and using q, the decoder recovers an approximation to x in the form of $\hat{x} := Fq$. The optimal choice of Q and F (dual to E) yields the analysis distortion. We will refer to this optimization problem as the "Analysis Problem". (Also see [18].)

There are some remarks to be made regarding the definition of the analysis distortion and its above interpretation. First, the assumption that the reconstruction must be "linear" is made for practical reasons. Non-linear methods (such as *consistent reconstruction* [21, 11]) have also been considered in the literature for the above type of encoders, but these methods generally come with higher computational cost. Second, the assumption that F be dual to E is made so that the reconstruction is guaranteed to be exact in the high-resolution quantization limit, or simply when there is no quantization. It should be noted that this assumption is also meaningful in order to make sense of how the choice of E might influence the reconstruction accuracy: without any assumptions on F or Q, F can be chosen optimally (in the sense of minimizing the synthesis distortion for $\mathcal{X}$) and Q can be chosen to produce the corresponding optimally quantized coefficients, thereby erasing the influence of E. To make this precise, let us define

$$\mathcal{D}(E, Q, F, \mathcal{X}) := \sup_{x \in \mathcal{X}} \|x - FQ(Ex)\|_2 \tag{4}$$

for any pair of frames (E, F) (not necessarily dual) and any quantization map $Q : \mathbb{R}^m \to \mathcal{A}^m$. The infimum in (1) is achieved for every x, so there is an optimal $q^*(x) \in \mathcal{A}^m$ that yields

$$\mathcal{D}_{\mathrm{s}}(F, \mathcal{A}, \mathcal{X}) = \sup_{x \in \mathcal{X}} \|x - Fq^*(x)\|_2.$$

Since E is injective, a map $Q^* : \mathbb{R}^m \to \mathcal{A}^m$ can be found such that $Q^*(Ex) = q^*(x)$, yielding

$$\mathcal{D}_{\mathrm{s}}(F, \mathcal{A}, \mathcal{X}) = \mathcal{D}(E, Q^*, F, \mathcal{X}) = \inf_{Q} \ \mathcal{D}(E, Q, F, \mathcal{X}).$$

This relation is valid for all E and F. Fixing E and minimizing over its dual frames F now results in the interpretation of the analysis distortion given in the previous paragraph:

$$\mathcal{D}_{\mathrm{a}}(E,\mathcal{A},\mathcal{X}) = \inf_{(Q,F):FE=I} \mathcal{D}(E,Q,F,\mathcal{X}). \tag{5}$$

Finally, it should be noted that in the setup of this paper we do not allow the decoder to pick a different linear map F depending on the quantized vector q that it receives. While such an adaptive method could be implemented in practice, it is not much different from (and sometimes equivalent to) using an arbitrary non-linear decoder.

The distortion performance of any given encoder-decoder scheme (adaptive or not) that uses at most N codewords for $\mathcal{X}$ is constrained by the entropic lower bound $\varepsilon(N,\mathcal{X})$ which is defined to be the smallest $\varepsilon > 0$ for which there exists an ε-net for $\mathcal{X}$ of cardinality N. Then we have

$$\mathcal{D}(E,Q,F,\mathcal{X}) \geq \mathcal{D}_{\mathrm{s}}(F,\mathcal{A},\mathcal{X}) \geq \mathcal{D}_{\mathrm{a}}(E,\mathcal{A},\mathcal{X}) \geq \varepsilon(|\mathcal{A}|^m,\mathcal{X}) \tag{6}$$

where the second inequality assumes that $FE = I$. The entropic lower bound depends on $\mathcal{A}$ only through its cardinality, so it will not necessarily be useful without a suitable matching between $\mathcal{A}$ and $\mathcal{X}$. For example, if the convex hull of $F(\mathcal{A}^m)$ is small compared to $\mathcal{X}$, then $\mathcal{D}_{\mathrm{s}}(F,\mathcal{A},\mathcal{X})$ will be artificially large due to empty regions in $\mathcal{X}$. For the rest of this paper, we will fix $\mathcal{X} = B_1$. With slight abuse of notation, define

$$\mathcal{D}_{\mathrm{s}}(F,L) := \inf_{\mathcal{A}:|\mathcal{A}|=L} \mathcal{D}_{\mathrm{s}}(F,\mathcal{A},B_1), \tag{7}$$

and

$$\mathcal{D}_{\mathrm{a}}(E,L) := \inf_{\mathcal{A}:|\mathcal{A}|=L} \mathcal{D}_{\mathrm{a}}(E,\mathcal{A},B_1). \tag{8}$$

With the simple volumetric bound $\varepsilon(N,B_1) \geq N^{-1/k}$, the relations in (6) yield

$$\mathcal{D}_{\mathrm{s}}(F,L) \geq \mathcal{D}_{\mathrm{a}}(E,L) \geq L^{-m/k}. \tag{9}$$

One of the main results of this paper is a matching upper bound for the analysis distortion for Gaussian random frames. We call E a *standard Gaussian random frame* if $(E_{i,j})$ are independent standard Gaussian random variables.

Theorem 1.1. *There exist absolute positive constants c_1 and c_2 such that for any scalar $\eta \in (0,1/2]$, if $k \geq \frac{8}{\eta}$ and $\frac{m}{k} \geq \frac{c_1}{\eta^2}\log\frac{m}{k}$, then with probability at least $1-\exp(-c_2\eta^2 m)$, a standard Gaussian random frame E of m vectors in $\mathbb{R}^k$ satisfies*

$$\mathcal{D}_{\mathrm{a}}(E,L) \leq \sqrt{k}L^{-(1-\eta)m/k}. \tag{10}$$

for all $L \geq 2$.

The proof of this theorem is constructive and algorithmic. Given η, we will, for each L, provide a quantization alphabet $\mathcal{A}$ which is an arithmetic progression of size L, and an algorithmic quantization map $Q : \mathbb{R}^m \to \mathcal{A}^m$ such that with high probability on E, there exists a certain dual frame F, determined by E and Q and computable by a fast, low-complexity algorithm, which yields $\mathcal{D}(E,Q,F,B_1) \lesssim (m/k)^{3/2}k^{1/2}L^{-(1-\eta)m/k}$. The chosen Q is a variation on vector-valued β-expansions and F is an alternative dual of E determined by Q. Suitably restricting the range of

m and k and summing over the failure probability for each L will then allow us to guarantee the distortion bound (10) uniformly over L.

Quantization of finite frames by sigma-delta modulation goes back to [2, 1]. Alternative duals for frames were introduced in the context of sigma-delta quantization in [16, 3], and more recently have been effectively used for Gaussian random frames in [12] and more generally for sub-Gaussian random frames in [15]. For these frames, an r-th order sigma-delta quantization scheme coupled with an associated Sobolev dual can achieve the bound $\mathcal{D}_{\mathrm{a}}(E, L) \lesssim L^{-1}(cr)^{2r}(m/k)^{-r}$ with high probability for some constant c. Optimizing this bound for r yields $\mathcal{D}_{\mathrm{a}}(E, L) \lesssim L^{-1}e^{-\tilde{c}\sqrt{m/k}}$ for another constant $\tilde{c}$ which is sub-optimal in terms of its dependence on both the alphabet size L and the oversampling factor m/k and L. In contrast, our result achieves near-optimal dependence on both parameters: Not only is the dependence on m/k exponential, but the rate can also be made arbitrarily close to its maximum value $\log_2 L$. To the best of our knowledge, this is the first result of this kind.[1]

The paper is organized as follows: In Section 2, we review the notion of general noise-shaping quantizers and their stability as well as alternative duals of frames based on noise shaping. In this section we introduce a number of new notions, namely, V-duals and V-condensation of frames and the concept of *distributed noise shaping*. In Section 3, we introduce the special case of *beta duals* and derive a general bound for the analysis distortion of an arbitrary frame using beta duals. As an example, we show in this section that beta duals are optimal for harmonic semicircle frames. In Section 4, we analyze the analysis distortion of Gaussian frames and prove Theorem 1.1 as well as some variations of it.

2 Quantization via distributed noise shaping

2.1 Noise-shaping quantizers

Let $\mathcal{A}$ be a quantization alphabet and J be a compact interval in $\mathbb{R}$. Let $h = (h_j)_{j\geq 0}$ be a given sequence, finite or infinite, and $h_0 = 1$. By a noise-shaping quantizer with the transfer sequence h, we mean any map $Q : J^m \to \mathcal{A}^m$ where $q := Q(y)$ satisfies

$$y - q = h * u$$

and $\|u\|_\infty \leq C$ for some constant C which is independent of m. Here $h * u$ refers to the convolution of h and u defined by $(h * u)_n := \sum h_j u_{n-j}$ where it is assumed that $u_n := 0$ for $n \leq 0$.

Noise-shaping refers to the fact that the "quantization noise" $y - q$ cannot be arbitrary. Even though the operator $H : u \mapsto h * u$ is invertible on $\mathbb{R}^m$; the requirement that $\|u\|_\infty$ must be controlled uniformly in m imposes restrictions on what q can be for a given y.

The above formulation amounts to saying that H is an $m \times m$ lower triangular Toeplitz matrix defined by h, in particular with unit diagonal. Let us relax the notion of a noise-shaping quantizer and assume that H is any given $m \times m$ lower triangular matrix with unit diagonal. We will refer

[1] In the case of conventional sigma-delta modulation, similar exponential error bounds that have been obtained previously are not compatible with the results of this paper: The results of [13] and [8] are for a fixed frame-like system, but using a different norm in infinite dimensions, and the dependence on L is unavailable. The results in [9], obtained in yet another functional setting for sigma-delta modulation, come close to being optimal, however these results were obtained under modeling assumptions on the quantization noise and circuit stability, and therefore are also not compatible with this paper.

to H as the noise transfer operator and keep the noise-shaping relation

$$y - q = Hu, \tag{11}$$

but since H is more general and m is therefore static, we can no longer insist that $\|u\|_\infty$ is controlled in the same manner. While we will mostly stick to convolution-based noise-shaping as defined at the beginning, the above relaxation is conceptually and algebraically simpler.

Noise-shaping quantizers do not exist unconditionally. However, under certain suitable assumptions on H and $\mathcal{A}$, they exist and can be implemented via recursive algorithms. The simplest is the greedy quantizer whose general formulation is given below:

Proposition 2.1. *Let $\mathcal{A} := \mathcal{A}_{L,\delta}$ denote the arithmetic progression in $\mathbb{R}$ which is of length L, spacing 2δ, and symmetric about 0. Assume that $H = I - \tilde{H}$, where $\tilde{H}$ is strictly lower triangular, and $\mu \geq 0$ such that $\|\tilde{H}\|_{\infty\to\infty} + \mu/\delta \leq L$. Then there exists a quantization map $Q_{H,\mathcal{A}} : [-\mu, \mu]^m \to \mathcal{A}^m$ implementable by a recursive algorithm such that $q := Q_{H,\mathcal{A}}(y)$ satisfies (11) and $\|u\|_\infty \leq \delta$.*

Proof. Since the endpoints of $\mathcal{A}$ are $\pm(L-1)\delta$, we have $|w - \mathrm{round}_{\mathcal{A}}(w)| \leq \delta$ for all $|w| \leq L\delta$. For $n = 1, \ldots, m$, we set

$$q_n := \mathrm{round}_{\mathcal{A}} \left(y_n + \sum_{j=1}^{n-1} \tilde{H}_{n,n-j} u_{n-j} \right)$$

and

$$u_n := y_n + \sum_{j=1}^{n-1} \tilde{H}_{n,n-j} u_{n-j} - q_n.$$

The proof now follows by induction. Setting $w = y_1$, we obtain $|u_1| \leq \delta$. Assuming that $|u_j| \leq \delta$ for $j = 1, \ldots, n-1$, we obtain

$$\left| y_n + \sum_{j=1}^{n-1} \tilde{H}_{n,n-j} u_{n-j} \right| \leq \mu + \delta \|\tilde{H}\|_{\infty\to\infty} \leq L\delta.$$

Hence $|u_n| \leq \delta$. □

Remark 2.2. *In the special case $Hu = h * u$ where $h_0 = 1$, note that $\|\tilde{H}\|_{\infty\to\infty} = \|h\|_1 - 1$.*

2.2 Alternative duals of frames for noise shaping

Let E be a frame and $y = Ex$ be the frame measurements of a given signal x. Assume that we quantize y using a noise-shaping quantizer with transfer operator H. Any left-inverse (dual) F of E gives

$$x - Fq = F(y - q) = FHu. \tag{12}$$

As explained in the introduction, the choice of F is to be made based upon E and the quantizer $Q_{H,\mathcal{A}}$, but not u which is unavailable to the decoder. We are interested in bounding $\|x - Fq\|_2$, but the best a priori bound for u is in the ∞-norm. This suggests that we employ the bound

$$\|x - Fq\|_2 \leq \|FH\|_{\infty\to 2} \|u\|_\infty. \tag{13}$$

With this bound, the natural objective would be to minimize $\|FH\|_{\infty\to 2}$ over all duals of E. However, an analytical solution for this problem is not available to us. If instead we minimize $\|FH\|_{2\to 2}$, then the solution is known explicitly, and given by $F = (H^{-1}E)^\dagger H^{-1}$, where $A^\dagger := (A^*A)^{-1}A^*$ is the pseudoinverse [3, 12].

More generally, for any $p \times m$ matrix V such that VE is also a frame for $\mathbb{R}^k$ (i.e. of rank k),

$$F_V := (VE)^\dagger V \tag{14}$$

is a dual of E, which we will call the *V-dual* of E. (For the case $V = H^{-1}$, the resulting dual was called the H-dual in [12] and denoted by F_H. With our more general definition, this dual more accurately becomes the H^{-1}-dual and denoted by $F_{H^{-1}}$.) The particular case $V = D^{-r}$, called the "Sobolev dual", was introduced in [3] and studied for random E in [12].

With a V-dual, we have $F_V H = (VE)^\dagger VH$. Given that $\|(VE)^\dagger\|_{2\to 2} = 1/\sigma_{\min}(VE)$, we employ the initial bound

$$\|F_V H\|_{\infty\to 2} \le \frac{\|VH\|_{\infty\to 2}}{\sigma_{\min}(VE)}.$$

Note that this bound is still valid when VE fails to be a frame, since then we have $\sigma_{\min}(VE) = 0$.

Let us compare the following two routes for bounding $\|VH\|_{\infty\to 2}$:

(i) $\|VH\|_{\infty\to 2} \le \sqrt{m}\|VH\|_{2\to 2}$. Since $\|VH\|_{2\to 2}/\sigma_{\min}(VE) \ge \|F_V H\|_{2\to 2} \ge 1/\sigma_{\min}(H^{-1}E)$, it follows that the minimum value of $\|VH\|_{2\to 2}/\sigma_{\min}(VE)$ is attained when $V = H^{-1}$. Hence, this route does not offer a new bound.

(ii) $\|VH\|_{\infty\to 2} \le \sqrt{p}\|VH\|_{\infty\to\infty}$. The minimization of $\sqrt{p}\|VH\|_{\infty\to\infty}/\sigma_{\min}(VE)$ opens up the possibility of a smaller error bound because $V = H^{-1}$ readily achieves the same bound as in (i).[2] Let $v_1, \dots, v_p$ denote the rows of V. Note that

$$\|VH\|_{\infty\to\infty} = \max_{1\le n\le p} \|v_n H\|_1.$$

Hence an effective strategy for designing V is based on choosing $v_1, \dots, v_p$ near (but not necessarily equal to) the bottom p left singular vectors of H. We will exploit this principle in the next section.

Note that different choices of V can yield the same V-dual. In particular, given any V, the $k\times m$ matrix $\tilde{V} := (W^*W)^{-1/2}W^*V$ yields $F_{\tilde{V}} = F_V$, where $W := VE$. This $\tilde{V}$ is not particularly useful for computational purposes, but it shows that p can be chosen as small as k without sacrificing performance. In general, when $p < m$, we will call VE the *V-condensation* of E.

2.3 Distributed noise shaping

By a *distributed* noise-shaping quantization scheme, we will mean a noise-shaping quantizer defined by a block-diagonal noise transfer matrix H, and for any given analysis frame E, an associated V-dual where V is also block-diagonal with matching sized blocks. More specifically, we set

$$H = \begin{bmatrix} H_1 & & \\ & \ddots & \\ & & H_l \end{bmatrix} \quad \text{and} \quad V = \begin{bmatrix} V_1 & & \\ & \ddots & \\ & & V_l \end{bmatrix}$$

[2] The following example shows that the minimum value can be strictly smaller: Let $H = \begin{bmatrix} 1 & 0 \\ -1 & 1 \end{bmatrix}$, $E = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$ for which $\sqrt{m}/\sigma_{\min}(H^{-1}E) = \sqrt{2/5}$. Meanwhile, $V = \begin{bmatrix} 1 & 1 \end{bmatrix}$ yields $\sqrt{p}\|VH\|_{\infty\to\infty}/\sigma_{\min}(VE) = 1/2$.

where $H_i \in \mathbb{R}^{m_i \times m_i}$ and $V_i \in \mathbb{R}^{p_i \times m_i}$ with $\sum_i m_i = m$ and $\sum_i p_i = p$. We further decompose E, y, q and u as

$$E = \begin{bmatrix} E_1 \\ \vdots \\ E_l \end{bmatrix}, \quad y = \begin{bmatrix} y_1 \\ \vdots \\ y_l \end{bmatrix}, \quad q = \begin{bmatrix} q_1 \\ \vdots \\ q_l \end{bmatrix} \quad \text{and} \quad u = \begin{bmatrix} u_1 \\ \vdots \\ u_l \end{bmatrix},$$

where $E_i \in \mathbb{R}^{m_i \times k}$, and y_i, q_i, $u_i \in \mathbb{R}^{m_i}$, $i \in [l] := \{1, \dots, l\}$. Note that the E_i may or may not be individual frames in $\mathbb{R}^k$. With this notation, we have l individual quantizers that run in parallel:

$$E_i x - q_i = H_i u_i, \qquad i \in [l]. \tag{15}$$

Let $v_{i,j}$ denote the jth row of V_i, $i \in [l]$, $j \in [p_i]$. Then we have

$$\|VH\|_{\infty \to \infty} = \max_{i \in [l]} \max_{j \in [p_i]} \|v_{i,j} H_i\|_1.$$

Hence, it follows that the reconstruction error obeys the bound

$$\|x - F_V q\|_2 \leq \frac{\sqrt{p}\|u\|_\infty}{\sigma_{\min}(VE)} \max_{i \in [l]} \max_{j \in [p_i]} \|v_{i,j} H_i\|_1.$$

Remark 2.3. *For conceptual clarity, we may assume that $l \neq 1$ so that distributed noise shaping is genuinely a special case of noise shaping as described in Section 2.2. However, we will not need to turn this into an explicit assumption because we will only consider $l \geq k$, and most often, $l > k$.*

Remark 2.4. *Our formalism for noise-shaping quantization and alternate-dual reconstruction assumes that the frame vectors and the measurements are processed in the given (specified) order. For distributed noise shaping, we assume that these vectors are collected into groups according to the same ordering.*

3 Beta duals

We now propose a specific distributed noise-shaping quantization scheme and corresponding alternative duals for reconstruction which will form the basis for the main results of this paper.

For any $\beta > 1$, let h^β be the (length-2) sequence given by $h_0^\beta = 1$ and $h_1^\beta = -\beta$.

Given $k \leq l \leq m$, let $\boldsymbol{m} := (m_1, \dots, m_l)$ and $\boldsymbol{\beta} := (\beta_1, \dots, \beta_l)$ be such that $\sum m_i = m$ and $\beta_i > 1$, $i \in [l]$. For each i, we set H_i equal to the convolution operator $u \mapsto h^{\beta_i} * u$, $u \in \mathbb{R}^{m_i}$, as defined in Section 2.1. Note that each H_i is an $m_i \times m_i$ bidiagonal matrix with diagonal entries equal to 1 and the subdiagonal entries equal to $-\beta_i$.

We set $p_i = 1$ and $V_i = v_{i,1} := [\beta_i^{-1} \;\; \beta_i^{-2} \;\; \cdots \;\; \beta_i^{-m_i}]$, $i \in [l]$. Let $V =: V_{\boldsymbol{\beta},\boldsymbol{m}}$ and H be as in Section 2.3. Note that $p = l$. It follows that $V_i H_i = [0 \;\; \cdots \;\; 0 \;\; \beta_i^{-m_i}]$, $\|V_i H_i\|_{\infty \to \infty} = \beta_i^{-m_i}$, and therefore $\|VH\|_{\infty \to \infty} = \beta_*^{-m_*}$, where $\beta_* := \min_i \beta_i$ and $m_* := \min_i m_i$. With the setup and the notation of Section 2.3, it follows that

$$\|x - F_{V_{\boldsymbol{\beta},\boldsymbol{m}}} q\|_2 \leq \frac{\sqrt{l}\|u\|_\infty}{\sigma_{\min}(V_{\boldsymbol{\beta},\boldsymbol{m}} E)} \beta_*^{-m_*}. \tag{16}$$

Note that this particular choice of $V_i = v_{i,1}$ is derived from β-expansions of real numbers. In fact, solving the difference equations (15) for q_i and recovering with $v_{i,1}$ gives a truncated β_i-expansion of $v_{i,1} E_i x$.

We will refer to $V_{\boldsymbol{\beta},\boldsymbol{m}}E$ as the $(\boldsymbol{\beta},\boldsymbol{m})$-*condensation* of E, and the corresponding $V_{\boldsymbol{\beta},\boldsymbol{m}}$-dual $F_{V_{\boldsymbol{\beta},\boldsymbol{m}}}$ as the $(\boldsymbol{\beta},\boldsymbol{m})$-*dual* of E. We will use the term *beta dual* when implicitly referring to any such dual.

In a typical application all the β_i would be set equal to some $\beta = \beta_*$ and the m_i (approximately) equal to m/l. To denote this choice, we will replace the vector pair $(\boldsymbol{\beta},\boldsymbol{m})$ by the scalar triplet (β, m, l), ignoring the exact values of the m_i as long as $m_* \geq \lfloor m/l \rfloor$.

3.1 Bounding the analysis distortion via beta duals

For any integer $L \geq 2$ and scalar $\mu > 0$, let

$$S_{\mu,L} := \{(\beta,\delta) : \beta > 1,\ \delta > 0,\ \beta + \mu/\delta \leq L\}. \tag{17}$$

This set corresponds to the admissible set of β and δ values such that the greedy quantization rule described in Proposition 2.1 using $\mathcal{A}_{L,\delta}$ and a noise shaping operator H as defined above (with all $\beta_i = \beta$) is stable with $\|u\|_\infty \leq \delta$ uniformly for all inputs bounded by μ. Note that the best ℓ_∞-norm bound for $y = Ex$ over all $x \in B_1$ is equal to $\|E\|_{2\to\infty}$ and that $S_{\mu,L}$ decreases with μ. Combining (16) with (4), (5) and (8), we immediately obtain the following result:

Proposition 3.1. *Let E be an $m \times k$ frame and $\mu \geq \|E\|_{2\to\infty}$. For any integer $L \geq 2$, we have*

$$\mathcal{D}_{\mathrm{a}}(E, L) \leq \inf_{\substack{(\beta,\delta)\in S_{\mu,L} \\ l \geq k}} \frac{\delta\sqrt{l}\beta^{-\lfloor m/l \rfloor}}{\sigma_{\min}(V_{\beta,m,l}E)}. \tag{18}$$

In order to find good upper bounds for $\mathcal{D}_{\mathrm{a}}(E, L)$ via this proposition, we will do the following:

(1) Find an upper bound μ for $\|E\|_{2\to\infty} = \max\limits_{1\leq i\leq m} \|e_i\|_2$.

(2) Given any $\beta > 1$ and $l \geq k$, find a lower bound on $\sigma_{\min}(V_{\beta,m,l}E)$.

(3) Find optimal (or near-optimal) values of β, δ, and l for the resulting effective upper bound replacing (18).

We will primarily carry out this procedure for random frames. However, before doing that in Section 4, we will discuss in Section 3.2 a specific family of (deterministic) frames, namely the harmonic semicircle frames in $\mathbb{R}^2$.

3.2 Example: the analysis distortion for the harmonic semicircle frames

Harmonic semicircle frames are obtained from harmonic frames by removing some of their symmetry (see [4] for the general definition). In $\mathbb{R}^2$, the harmonic semicircle frame of size m is defined by the $m \times 2$ matrix $E := E_{\mathrm{hsc},m}$ given by

$$E_{\mathrm{hsc},m} := \begin{bmatrix} \cos \pi/m & \sin \pi/m \\ \cos 2\pi/m & \sin 2\pi/m \\ \vdots & \vdots \\ \cos \pi & \sin \pi \end{bmatrix}$$

We set $l = 2$, split E into E_1 and E_2 as in Section 2.3, and for simplicity, assume that m is even. (The odd case can be handled with a slight alteration.)

For any $\beta > 1$, consider $V_{\beta,m,2}E$, which is given by the 2×2 matrix

$$\begin{bmatrix} v_\beta E_1 \\ v_\beta E_2 \end{bmatrix}$$

where $v_\beta := V_1 = V_2 = [\beta^{-1} \cdots \beta^{-m/2}]$. Note that each row of E_2 is a 90 degree rotation of the corresponding row of E_1. Therefore $v_\beta E_2$ is also a 90 degree rotation of $v_\beta E_1$ which implies that $\sigma_{\min}(V_{\beta,m,2}E) = \|v_\beta E_1\|_2 =: C_{\beta,m}$. An exact expression for $C_{\beta,m}$ can be found easily, namely,

$$C_{\beta,m} = \frac{(1+\beta^{-m})^{1/2}}{|\beta - e^{\pi i/m}|}.$$

It is easily seen that $C_{\beta,m} \geq \beta^{-1}$ for all m (with equality if and only if $m = 2$).

Since all frame vectors are of unit norm, we can set $\mu = 1$. Finally, we choose $\beta = mL/(1+m)$ and $\delta = (1+m)/L$, which are near optimal values as suggested by Proposition 3.1. With these values, it follows that

$$\mathcal{D}_{\mathrm{a}}(E_{\mathrm{hsc},m}, L) < \sqrt{2e} m L^{-m/2}.$$

In Figure 1 we illustrate how the $(\beta, m, 2)$-dual for $m = 12$ and $\beta = 1.6$.

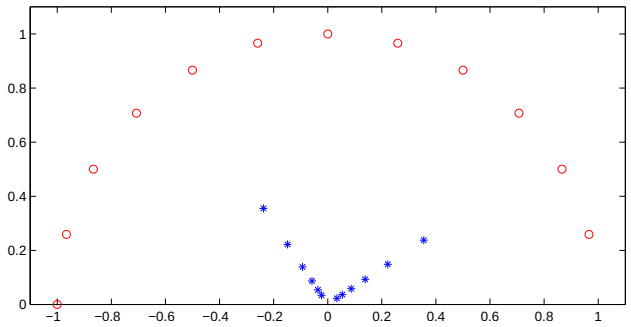

Figure 1: Harmonic semicircle frame using circles and (β, m, l)-dual frame with stars, $\beta = 1.6$, $m = 12$, $l = 2$.

4 Analysis distortion of Gaussian frames

In this section, we assume that E is an $m \times k$ random matrix with entries drawn independently from $\mathcal{N}(0,1)$. As before, L stands for a given alphabet size. To bound the analysis distortion $\mathcal{D}_{\mathrm{a}}(E, L)$, we will again follow the three step plan outlined in section 3.1 with a slight modification: Note that the entries of the $l \times k$ matrix $V_{\beta,m,l}E$ are mutually independent and the entries of the ith row are identically distributed according to $\mathcal{N}(0, \sigma^2_{m_i})$, where

$$\sigma_n^2 := \sum_{j=1}^{n} \beta^{-2j}$$

and the m_i satisfy $\sum m_i = m$ and $m_i \geq \lfloor m/l \rfloor$ for all i. Our analysis of $\sigma_{\min}(V_{\beta,m,l}E)$ would be simpler if all the σ_{m_i} were equal, which holds if and only if m is a multiple of l. Since l is a parameter which will be optimized later, this condition cannot be enforced in advance. To circumvent this issue, we define for each l, the submatrix $\tilde{E}_l$ of E formed by the first $\tilde{m}$ rows of E where $\tilde{m} := \lfloor m/l \rfloor l$. Provided that $\tilde{E}_l$ is a frame for $\mathbb{R}^k$ (which holds almost surely if $l \geq k$), the definition of the analysis distortion $\mathcal{D}_{\mathrm{a}}(E, \mathcal{A}, \mathcal{X})$ given by (3) implies that

$$\mathcal{D}_{\mathrm{a}}(E, \mathcal{A}, \mathcal{X}) \leq \mathcal{D}_{\mathrm{a}}(\tilde{E}_l, \mathcal{A}, \mathcal{X}),$$

because for any $\tilde{F}$ such that $\tilde{F}\tilde{E}_l = I$, the zero-padded $k \times m$ matrix $F := [\tilde{F}\ 0]$ satisfies $FE = I$ and

$$\mathcal{D}_{\mathrm{s}}(F, \mathcal{A}, \mathcal{X}) = \mathcal{D}_{\mathrm{s}}(\tilde{F}, \mathcal{A}, \mathcal{X}).$$

Hence we replace (18) with

$$\mathcal{D}_{\mathrm{a}}(E, L) \leq \inf_{l \geq k} \quad \mathcal{D}_{\mathrm{a}}(\tilde{E}_l, L) \tag{19}$$

where

$$\mathcal{D}_{\mathrm{a}}(\tilde{E}_l, L) \leq \inf_{(\beta,\delta) \in S_{\mu,L}} \quad \frac{\delta\sqrt{l}\beta^{-\lfloor m/l \rfloor}}{\sigma_{\min}(V_{\beta,\tilde{m},l}\tilde{E}_l)}. \tag{20}$$

4.1 Upper bound for $\|E\|_{2\to\infty}$

Note that because the Gaussian distribution has unbounded support, there is no absolute bound on $\|E\|_{2\to\infty}$ that is applicable almost surely, so we employ concentration bounds which hold with high probability. Let

$$\mathcal{M} := \left\{ \|E\|_{2\to\infty} \leq \mu := 4\sqrt{m} \right\} \tag{21}$$

The following well-known result will be suitable for this purpose:

Proposition 4.1 ([6, Theorem II.13][7])**.** *Let $m \geq k$ and E be an $m \times k$ random matrix whose entries are i.i.d. standard Gaussian variables. For any $t > 0$,*

$$\mathbb{P}\left(\left\{ \sigma_{\max}(E) \geq \sqrt{m} + \sqrt{k} + t \right\} \right) < e^{-t^2/2}.$$

Setting $t = 2\sqrt{m}$ and noting that $\sigma_{\max}(E) = \|E\|_{2\to 2} \geq \|E\|_{2\to\infty}$, we have

$$\mathbb{P}(\mathcal{M}^c) \leq \mathbb{P}\left(\left\{ \|E\|_{2\to 2} > 4\sqrt{m} \right\} \right) < e^{-2m}. \tag{22}$$

4.2 Lower bound for $\sigma_{\min}(V_{\beta,\tilde{m},l}\tilde{E}_l)$

Again, we employ concentration bounds. For any $\gamma > 0$, we are interested in the event

$$\mathcal{E}(\gamma) := \mathcal{E}(\gamma, \beta, l) := \left\{ \sigma_{\min}(V_{\beta,\tilde{m},l}\tilde{E}_l) \geq \gamma \right\}. \tag{23}$$

Note that the entries of the $l \times k$ matrix $V_{\beta,\tilde{m},l}\tilde{E}_l$ are now independent and identically distributed, with each entry having distribution $\mathcal{N}(0, \sigma^2_{\lfloor m/l \rfloor})$. It will suffice to use standard methods and results on the concentration of the minimum singular value of Gaussian matrices. The nature of these results differ depending on whether $l = k$ or $l > k$ and the optimal value of l depends on the desired level of probabilistic assurance.

For $l = k$, we use the following result:

Proposition 4.2 ([20, Theorem 3.1], [10])**.** *Let Ω be a $k \times k$ random matrix with entries drawn independently from $\mathcal{N}(0,1)$. Then for any $\varepsilon > 0$,*

$$\mathbb{P}\left(\left\{\sigma_{\min}(\Omega) \leq \frac{\varepsilon}{\sqrt{k}}\right\}\right) \leq \varepsilon.$$

For $l > k$, there are standard estimates concentration estimates (e.g. [6, 19, 22]) which take the lower end of the spectrum $\sqrt{l} - \sqrt{k}$ as a reference point. In contrast, the following result concerns the probability of the smallest singular value being near zero.

Theorem 4.3. *Let $l > k$ and Ω be an $l \times k$ random matrix whose entries are drawn independently from $\mathcal{N}(0,1)$. Then for any $0 < \varepsilon < 1$,*

$$\mathbb{P}\left(\left\{\sigma_{\min}(\Omega) \leq \varepsilon\sqrt{l}/2\right\}\right) \leq \left(10 + 8\sqrt{\log \varepsilon^{-1}}\right)^k e^{l/2}\varepsilon^{l-k}.$$

For a proof, see Appendix A. Note that there exists $\varepsilon_0 := \varepsilon_0(k,l) < 1$ such that this probability bound is not viable (i.e., exceeds 1) for $\varepsilon > \varepsilon_0$. This may turn out be the case in Theorem 4.6 (see Section 4.3.2 below) if neither m/l nor L is sufficiently large. However, the stronger assumptions in Corollary 4.7 and Theorem 1.1 rule out this possibility.

4.3 Bounding the analysis distortion via optimal choice of β and δ

4.3.1 Choosing $l = k$

Let $\eta \in (0,1)$ be arbitrary. We employ Proposition 4.2 with $\varepsilon = \beta^{-\eta\lfloor m/k \rfloor}$ which corresponds to $\gamma = \frac{\varepsilon}{\sqrt{k}}\sigma_{\lfloor m/k \rfloor}$ in (23). With the simple bound $\sigma_{\lfloor m/k \rfloor} \geq \beta^{-1} \geq L^{-1}$, the analysis distortion bound (20) yields

$$\mathcal{D}_{\mathrm{a}}(\tilde{E}_k, L) \leq \inf_{(\beta,\delta)\in S_{\mu,L}} kL\delta\beta^{-(1-\eta)\lfloor m/k \rfloor} \tag{24}$$

for all E in the event $\mathcal{E}(\gamma) \cap \mathcal{M}$, which occurs with probability at least $1 - \beta^{-\eta\lfloor m/k \rfloor} - e^{-2m}$. Note that $\mu = 4\sqrt{m}$.

A straightforward calculation shows that for any $\alpha > 1/(L-1)$, the minimum of $\delta\beta^{-\alpha}$ over the set $S_{\mu,L}$ is equal to $\mu\alpha\left(\frac{1+\alpha}{\alpha L}\right)^{1+\alpha}$, which is attained at $\beta = L\alpha/(1+\alpha) > 1$ and $\delta = \mu(1+\alpha)/L$.

We use these values of β and δ in the bound (24) with $\alpha = (1-\eta)\lfloor m/k \rfloor$ and $\mu = 4\sqrt{m}$. Noting that $(1 + 1/\alpha)^\alpha < e$, we have proved

Theorem 4.4. *Let $m \geq k$ and $L \geq 2$, except assume $L > 2$ if $m < 2k$. Suppose E is an $m \times k$ random matrix with entries drawn independently from $\mathcal{N}(0,1)$. For any $0 < \eta < 1 - \frac{1}{\lfloor m/k \rfloor (L-1)}$, the analysis distortion for E satisfies the bound*

$$\mathcal{D}_{\mathrm{a}}(E, L) < 8em^{3/2}L^{-(1-\eta)\lfloor m/k \rfloor}$$

with probability at least $1 - L^{-\eta\lfloor m/k \rfloor}e^{\eta/(1-\eta)} - e^{-2m}$.

Remark 4.5. *The most favorable choice of η depends on m, k, L as well as the probabilistic assurance that is desired. For example, as $\lambda := m/k \to \infty$, $\eta = \lambda^{-1/2}$ yields $\mathcal{D}_{\mathrm{a}}(E, L) = O(m^{3/2}L^{-\lambda(1+o(1))})$ with probability $1 - o(1)$.*

4.3.2 Choosing $l > k$

Again, let $\eta \in (0,1)$. We now employ Theorem 4.3 with $\varepsilon = \beta^{-\eta\lfloor m/l\rfloor}$ which corresponds to $\gamma = \varepsilon\sigma_{\lfloor m/l\rfloor}\sqrt{l}/2$. Again using $\sigma_{\lfloor m/l\rfloor} \geq L^{-1}$, the analysis distortion bound (20) becomes

$$\mathcal{D}_{\mathrm{a}}(\tilde{E}_l, L) \leq \inf_{(\beta,\delta)\in S_{\mu,L}} 2L\delta\beta^{-(1-\eta)\lfloor m/l\rfloor} \tag{25}$$

for all E in the event $\mathcal{E}(\gamma,\beta,l)\cap\mathcal{M}$, which occurs with probability at least

$$1 - e^{l/2}(6 + 4\sqrt{\eta\lfloor m/l\rfloor\log\beta})^k\beta^{-\eta\lfloor m/l\rfloor(l-k)} - e^{-2m}.$$

Motivated by the same calculation as in the $l = k$ case, we choose $\alpha = (1-\eta)\lfloor m/l\rfloor$, $\beta = L\alpha/(1+\alpha)$ and $\delta = \mu(1+\alpha)/L$, which now shows:

Theorem 4.6. *Let $m \geq l > k$ and $L \geq 2$, except assume $L > 2$ if $m < 2l$. Suppose E is an $m\times k$ random matrix with entries drawn independently from $\mathcal{N}(0,1)$. For any $0 < \eta < 1 - \frac{1}{\lfloor m/l\rfloor(L-1)}$, the analysis distortion for E satisfies the bound*

$$\mathcal{D}_{\mathrm{a}}(E, L) < 16em^{3/2}l^{-1}L^{-(1-\eta)\lfloor m/l\rfloor} \tag{26}$$

with probability at least

$$1 - e^{l/2}(10 + 8\sqrt{\eta\lfloor m/l\rfloor\log L})^k e^{\frac{\eta}{1-\eta}(l-k)}L^{-\eta\lfloor m/l\rfloor(l-k)} - e^{-2m}.$$

In particular, this result implies that the failure probability can be arranged to be exponentially small in m while ensuring that the distortion bound is still of the same order as in Theorem 4.4. Indeed, let $l = k + \lfloor\theta k\rfloor$ for some $\theta \geq 1/k$ (so that $l > k$). With this value of l, we have

$$\mathcal{D}_{\mathrm{a}}(E, L) < 16em^{3/2}k^{-1}L^{-(1-\eta)\lfloor\frac{m}{(1+\theta)k}\rfloor} \tag{27}$$

for all E in the event $\mathcal{E}(\gamma,\beta,l)\cap\mathcal{M}$, where the probability of $\mathcal{E}(\gamma,\beta,l)^c$ is bounded by

$$\exp\left\{k\left(\frac{1}{2}(1+\theta) + c_0 + \frac{1}{2}\log\frac{m}{k} + \frac{1}{2}\log\log L + \theta\frac{\eta}{1-\eta}\right) - \eta\left\lfloor\frac{m}{(1+\theta)k}\right\rfloor\lfloor\theta k\rfloor\log L\right\}$$

with $c_0 := \log(21)$. (In deriving this simplification, we have used that $k < l \leq (1+\theta)k$ and $10 + 8\sqrt{\eta\lfloor m/l\rfloor\log L} \leq 21\sqrt{(m/k)\log L}$.) If $k \geq 8$, $2/k \leq \theta \leq 1/4$ and $m/k \geq 5/2$, we have

$$\left\lfloor\frac{m}{(1+\theta)k}\right\rfloor\lfloor\theta k\rfloor \geq \frac{m\theta}{4(1+\theta)} \geq \frac{m\theta}{5}$$

so that the probability of $\mathcal{E}(\gamma,\beta,l)^c$ is further bounded by

$$\exp\left\{k\left(\frac{1}{2}(1+\theta) + c_0 + \frac{1}{2}\log\frac{m}{k} + \frac{1}{2}\log\log L + \theta\frac{\eta}{1-\eta}\right) - \frac{1}{5}\eta\theta m\log L\right\}.$$

Let $\theta \in (0, 1/4]$ and $\eta \in (0, 1/2]$ be fixed. There exists an absolute constant $c_1 > 0$ such that if $\frac{m}{k} \geq \frac{c_1}{\eta\theta}\log\frac{m}{k}$ and $k \geq 2/\theta$, we have

$$k\left(\frac{1}{2}(1+\theta) + c_0 + \frac{1}{2}\log\frac{m}{k} + \frac{1}{2}\log\log L + \theta\frac{\eta}{1-\eta}\right) \leq \frac{1}{10}\eta\theta m\log L$$

so that

$$\mathbb{P}(\mathcal{E}(\gamma,\beta,l)^c) \leq L^{-\frac{1}{10}\eta\theta m}. \tag{28}$$

(It can be checked that a sufficiently large value of c_1 implies a correspondingly large lower bound on m/k, which in turn guarantees that $1 - \frac{1}{\lfloor m/l \rfloor (L-1)} > \eta$.) Note that for fixed values of η and θ, the event $\mathcal{E}(\gamma,\beta,l)$ depends on L through γ and β. Since $m \geq 20/\eta\theta$ (by readjusting the value of c_1 if necessary), it follows that

$$\mathbb{P}\left(\bigcup_{L=2}^{\infty} \mathcal{E}(\gamma,\beta,l)^c\right) \leq \sum_{L=2}^{\infty} L^{-\frac{1}{10}\eta\theta m} \leq 3 \cdot 2^{-\frac{1}{10}\eta\theta m}. \tag{29}$$

Note also that if $m/k \geq (1+\theta)/\theta^2$, then

$$\left\lfloor \frac{m}{(1+\theta)k} \right\rfloor \geq \frac{m}{(1+\theta)k} - 1 \geq (1-\theta)\frac{m}{k}$$

so that the error bound (27) can be simplified to

$$\mathcal{D}_{\mathrm{a}}(E,L) < 16em^{3/2}k^{-1}L^{-(1-\eta-\theta)\frac{m}{k}}.$$

Choosing $\theta = \eta \in (0, 1/4)$, replacing η by $\eta/2$, we obtain the following corollary:

Corollary 4.7. *There exists an absolute constant $c_1 > 0$ such that if $\eta \in (0,1/2)$, and m and k are such that $k \geq 4/\eta$ and $\frac{m}{k} \geq \frac{c_1}{\eta^2}\log\frac{m}{k}$, an $m \times k$ random matrix E with entries drawn independently from $\mathcal{N}(0,1)$ satisfies*

$$\mathcal{D}_{\mathrm{a}}(E,L) < 16em^{3/2}k^{-1}L^{-(1-\eta)\frac{m}{k}}$$

for all $L \geq 2$ with probability at least $1 - 3 \cdot 2^{-\frac{1}{40}\eta^2 m} - e^{-2m}$.

Proof of Theorem 1.1. We make the following observations. First, increasing the value of c_1 in Corollary 4.7 if necessary, the assumption $\frac{m}{k} \geq \frac{c_1}{\eta^2}\log\frac{m}{k}$ implies that $\frac{m}{k} \geq \frac{c_1}{\eta^2}$, and therefore

$$16e\left(\frac{m}{k}\right)^{3/2} < e^{5+1.5\log(m/k)} < L^{c_3\eta^2 m/k},$$

where $c_3 \leq 13/c_1$. Again increasing the value of c_1 if necessary, we can assume that $c_3 \leq 4$ so that $1 - \eta - c_3\eta^2 \geq 1 - 2\eta$. Next, we can restrict $\eta \in (0, 1/4]$, replace η by $\eta/2$ as before, and replace c_1 with $4c_1$. This yields the error bound given in (10). For the probability guarantee, note that $m \geq c_1/\eta^3$. Once again increasing the value of c_1 if necessary, we can find an absolute constant $c_2 > 0$ such that $3 \cdot 2^{-\frac{1}{160}\eta^2 m} + e^{-2m} \leq e^{-c_2\eta^2 m}$ for all admissible m, hence Theorem 1.1 follows.

Acknowledgements

The authors would like thank Thao Nguyen, Rayan Saab and Özgür Yılmaz for useful conversations.

A Proof of Theorem 4.3

The proof of Theorem 4.3 will use a standard covering argument. Note that for each fixed $z \in \mathbb{R}^k$ with unit norm, Ωz has entries that are i.i.d. $\mathcal{N}(0,1)$. For any such vector we have the following basic lemma:

Lemma A.1. *Let $\xi \sim \mathcal{N}(0, I_l)$. For any $0 < \varepsilon \leq 1$, we have*

$$\mathbb{P}\left(\left\{\|\xi\|_2 \leq \varepsilon\sqrt{l}\right\}\right) \leq \varepsilon^l e^{(1-\varepsilon^2)l/2}. \tag{30}$$

Proof. For any $t \geq 0$, we have

$$\mathbb{P}\left(\left\{\|\xi\|_2^2 \leq \varepsilon^2 l\right\}\right) \leq \int_{\mathbb{R}^l} e^{(\varepsilon^2 l - \|x\|_2^2)t/2} \frac{e^{-\|x\|_2^2/2}}{(2\pi)^{l/2}}\,\mathrm{d}x = e^{\varepsilon^2 tl/2} \int_{\mathbb{R}^l} \frac{e^{-(1+t)\|x\|_2^2/2}}{(2\pi)^{l/2}}\,\mathrm{d}x = \left(\frac{e^{\varepsilon^2 t}}{1+t}\right)^{l/2}.$$

Choosing $t = \varepsilon^{-2} - 1$ yields the desired bound. □

For an arbitrary $\tau > 1$, let $\mathcal{E}_1$ be the event $\{\|\Omega\|_{2\to 2} \leq 2\tau\sqrt{l}\}$. Proposition 4.1 with $m = l$, $E = \Omega$, $t = 2(\tau-1)\sqrt{l}$ implies

$$\mathbb{P}(\mathcal{E}_1^c) \leq e^{-2(\tau-1)^2 l}.$$

Next, consider a ρ-net Q of the unit sphere of $\mathbb{R}^k$ with $|Q| \leq (1 + 2/\rho)^k$, where $\rho = \varepsilon/(4\tau)$. Let $\mathcal{E}_2$ be the event $\{\|\Omega z\|_2 \geq \varepsilon\sqrt{l},\ \forall z \in Q\}$. By Lemma A.1, we have

$$\mathbb{P}(\mathcal{E}_2^c) \leq |Q|\varepsilon^l e^{(1-\varepsilon^2)l/2} \leq (\varepsilon + 8\tau)^k \varepsilon^{l-k} e^{(1-\varepsilon^2)l/2}.$$

Let Ω be chosen in the event $\mathcal{E}_1 \cap \mathcal{E}_2$. For any unit norm $x \in \mathbb{R}^k$, there exists a $z \in Q$ such that $\|x - z\|_2 \leq \rho$. Then $\|\Omega(x-z)\|_2 \leq 2\tau\rho\sqrt{l} = \varepsilon\sqrt{l}/2$ and $\|\Omega z\|_2 \geq \varepsilon\sqrt{l}$, so that

$$\|\Omega x\|_2 \geq \|\Omega z\|_2 - \|\Omega(x-z)\|_2 \geq \varepsilon\sqrt{l}/2,$$

hence $\sigma_{\min}(\Omega) \geq \varepsilon\sqrt{l}/2$. It follows that $\mathcal{F} := \left\{\sigma_{\min}(\Omega) \leq \varepsilon\sqrt{l}/2\right\} \subset \mathcal{E}_1^c \cup \mathcal{E}_2^c$, and therefore

$$\mathbb{P}(\mathcal{F}) \leq e^{-2(\tau-1)^2 l} + (\varepsilon + 8\tau)^k \varepsilon^{l-k} e^{l/2}.$$

We still have the freedom to choose $\tau > 1$ as a function of ε, l, and k. For simplicity, we choose $\tau = 1 + \sqrt{\log \varepsilon^{-1}}$ so that $e^{-2(\tau-1)^2 l} = \varepsilon^{2l}$. Noting that $1 + (1+8\tau)^k \leq (2+8\tau)^k$, we obtain

$$\mathbb{P}(\mathcal{F}) < \varepsilon^{l-k}\left(1 + (1+8\tau)^k e^{l/2}\right) < \left(10 + 8\sqrt{\log \varepsilon^{-1}}\right)^k e^{l/2}\varepsilon^{l-k}.$$

References

[1] J.J. Benedetto, A.M. Powell, and Ö. Yılmaz. Second-order sigma-delta ($\Sigma\Delta$) quantization of finite frame expansions. *Appl. Comput. Harmon. Anal.*, 20(1):126–148, 2006.

[2] J.J. Benedetto, A.M. Powell, and Ö. Yılmaz. Sigma-delta quantization and finite frames. *Information Theory, IEEE Transactions on*, 52(5):1990–2005, 2006.

[3] J. Blum, M. Lammers, A.M. Powell, and Ö. Yılmaz. Sobolev duals in frame theory and sigma-delta quantization. *Journal of Fourier Analysis and Applications*, 16(3):365–381, 2010.

[4] B.G. Bodmann and V.I. Paulsen. Frame paths and error bounds for sigma-delta quantization. *Appl. Comput. Harmon. Anal.*, 22(2):176–197, 2007.

[5] B. Chazelle. *The discrepancy method: randomness and complexity*. Cambridge University Press, 2002.

[6] K.R. Davidson and S.J. Szarek. Local operator theory, random matrices and Banach spaces. In *Handbook of the geometry of Banach spaces, Vol. I*, pages 317–366. North-Holland, Amsterdam, 2001.

[7] K.R. Davidson and S.J. Szarek. Addenda and corrigenda to: "Local operator theory, random matrices and Banach spaces". In *Handbook of the geometry of Banach spaces, Vol. 2*, pages 1819–1820. North-Holland, Amsterdam, 2003.

[8] P. Deift, F. Krahmer, and C.S. Güntürk. An optimal family of exponentially accurate one-bit sigma-delta quantization schemes. *Communications on Pure and Applied Mathematics*, 64(7):883–919, 2011.

[9] M.S. Derpich, E.I. Silva, D.E. Quevedo, and G.C. Goodwin. On optimal perfect reconstruction feedback quantizers. *IEEE Trans. Signal Process.*, 56(8, part 2):3871–3890, 2008.

[10] A. Edelman. Eigenvalues and condition numbers of random matrices. *SIAM Journal on Matrix Analysis and Applications*, 9(4):543–560, 1988.

[11] V.K. Goyal, M. Vetterli, and N.T. Thao. Quantized overcomplete expansions in $\mathbb{R}^N$: analysis, synthesis, and algorithms. *Information Theory, IEEE Transactions on*, 44(1):16–31, 1998.

[12] C. S. Güntürk, M. Lammers, A. M. Powell, R. Saab, and Ö. Yılmaz. Sobolev duals for random frames and $\Sigma\Delta$ quantization of compressed sensing measurements. *Found. Comput. Math.*, 13(1):1–36, 2013.

[13] C.S. Güntürk. One-bit sigma-delta quantization with exponential accuracy. *Communications on Pure and Applied Mathematics*, 56(11):1608–1630, 2003.

[14] C.S. Güntürk. Mathematics of analog-to-digital conversion. *Comm. Pure Appl. Math.*, 65(12):1671–1696, 2012.

[15] F. Krahmer, R. Saab, and Ö. Yılmaz. Sigma-delta quantization of sub-gaussian frame expansions and its application to compressed sensing. *Information and Inference*, 3(1):40–58, 2014.

[16] M. Lammers, A.M. Powell, and Ö. Yılmaz. Alternative dual frames for digital-to-analog conversion in sigma-delta quantization. *Adv. Comput. Math.*, 32(1):73–102, 2010.

[17] J. Matoušek. *Geometric discrepancy*, volume 18 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 1999. An illustrated guide.

[18] A.M. Powell, R. Saab, and Ö. Yılmaz. Quantization and finite frames. In *Finite Frames*, pages 267–302. Springer, 2013.

[19] M. Rudelson and R. Vershynin. Smallest singular value of a random rectangular matrix. *Communications on Pure and Applied Mathematics*, 62(12):1707–1739, 2009.

[20] M. Rudelson and R. Vershynin. Non-asymptotic theory of random matrices: extreme singular values. In *Proceedings of the International Congress of Mathematicians. Volume III*, pages 1576–1602. Hindustan Book Agency, New Delhi, 2010.

[21] N.T. Thao and M. Vetterli. Lower bound on the mean-squared error in oversampled quantization of periodic signals using vector quantization analysis. *Information Theory, IEEE Transactions on*, 42(2):469–479, 1996.

[22] R. Vershynin. Introduction to the non-asymptotic analysis of random matrices. In *Compressed sensing*, pages 210–268. Cambridge Univ. Press, Cambridge, 2012.